

Extreme Data Science

Sudip Dosanjh, Shane Canon, Jack DeSlippe, Kjiersten Fagnan,
Richard Gerber, Lisa Gerhardt, Jason Hick, Douglas Jacobsen,
David Skinner, and Nicholas J. Wright

Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720

Exascale computing has received considerable attention in the high performance computing (HPC) community and commonly discussed drivers include predictive computer simulations and uncertainty quantification. Extreme data science is very important to the U.S. Department of Energy as its experimental facilities become increasingly inundated with large quantities of scientific data. Sometimes exascale computing is viewed as being at odds with the needs of big data. However, many of the underlying technical challenges (energy efficiency, concurrency, memory, reliability, programmability) are common. It is evident that extreme data science also drives the need for exascale computing technologies.

Over the past four years, more data has been transferred to NERSC than away from NERSC – an unprecedented paradigm shift for a supercomputing center. This trend is driven by exponential increases in detector resolution as well as increased automation and high-throughput techniques. Data pipelines to NERSC computers increased in bandwidth 10x to 100GB based networking. More than a petabyte of biosciences, climate and high energy physics data is transferred to NERSC every month. This data explosion is expected to continue as data generation rates at experimental facilities increase and our ability to move and process data improves. New discoveries based on this data most often require a mixture of data logistics, data management and computing-driven analysis capabilities that ideally would be found in a single scalable scientific computing architecture.

Data science capabilities are playing an increasingly central role in major scientific discoveries. Data resources at NERSC such as parallel filesystems, data transfer nodes, databases and science gateways now play a major role recent high impact science: 1) the measurement of the Θ_{13} neutrino mixing angle, *Science Magazine's* top ten breakthroughs of 2012, 2) the international Planck collaboration's creation of the most detailed maps yet of the cosmic microwave background, 3) the data-driven discovery of the youngest supernovae ever observed (PTF11kly) and others all make heavy use of data science capabilities deployed at NERSC. The rising tide of data science is informed by earlier work in this area. Since 1996, the high energy and nuclear physics communities and NERSC have partnered to deploy and maintain a large cluster, called PDSF (Parallel Distributed Systems Facility), to help scientists analyze large data sets. Expanding this model has been a success, in 2010 the Joint Genome Institute (www.jgi.doe.gov) and NERSC entered into a partnership

to deploy the Genepool data cluster, enabling researchers to sequence and analyze over 50 terabases (approximately 20,000 human genomes) this year alone.

We draw on these experiences to inform future strategies for extreme scale data science in the context of HPC centers such as NERSC. Systems architected for either compute-heavy or data-intensive workloads are derived largely from the same technological components. Careful attention to balance between these needs and constraints can enable a class of HPC resources that are likewise capable as extreme scale data science resources.

NERSC's most recent supercomputer is Edison, a 5200 node Cray XC30 Cascade machine was deployed in 2013. Instead of focusing on floating point operations, Edison was designed to optimize data motion, which is critical to both predictive simulation and data-intensive computing. The memory capacity, memory bandwidth, bisection bandwidth and I/O bandwidth have been architected to achieve balance between compute- and data-intensive demands. As scientific computing architecture moves toward exascale, a question that is the subject of research at NERSC and numerous other facilities is how to meet both computing and data demands in the most effective way. So far, our investigations indicate that we expect some differences between compute-intensive and data-intensive focused exascale machines. We expect a compute-intensive architecture to maximize computational density and local bandwidth for a given power/cost constraint, whereas a data-intensive architecture will maximize data capacity and global bandwidth for a given power/cost constraint. It will also be critical to develop new algorithms and software for analyzing and sharing extremely large science data sets. Successfully architected balance between compute and data has significant upside for the scientific computing community allowing for flexible *in-situ* adaptation between compute and data-driven tasks. The Edison system at NERSC hope to inform extreme scale data science in finding computing and data analysis solutions

Acknowledgement

This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.